

WHAT IS CLAIMED IS:

1. A cluster comprising a plurality of computer systems, wherein each of the plurality of computer systems is configured to execute one or more virtual machines, each of the plurality of computer systems comprising hardware and a plurality of instructions that, when executed on the hardware, detects that a first load of a first computer system of the plurality of computer systems exceeds a second load of a second computer system of the plurality of computer systems and migrates at least a first virtual machine executing on the first computer system to the second computer system responsive to detecting that the first load exceeds the second load.
10
2. The cluster as recited in claim 1 wherein the first virtual machine executes on the second computer system independent of the first computer system, even if the first virtual machine was initially launched on the first computer system.
15
3. The cluster as recited in claim 1 wherein the plurality of instructions, when executed on the first computer system, select the second computer system to compare loads.
4. The cluster as recited in claim 3 wherein the plurality of instructions, when executed on the first computer system, randomly select the second computer system from the plurality of computer systems.
20
5. The cluster as recited in claim 1 wherein the first virtual machine has a corresponding load that is nearest, among loads of the virtual machines executing on the first computer system, to 1/2 the difference between the first load and the second load.
25
6. The cluster as recited in claim 5 wherein the corresponding load of the first virtual machine represents the actual load experienced in executing the first virtual machine on the first computer system.

7. The cluster as recited in claim 5 wherein the corresponding load of the first virtual machine represents the target load programmed for the first virtual machine in the first computer system.

5

8. The cluster as recited in claim 1 wherein the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load, wherein the first computer system is configured to transmit one or 10 more load factors to the second computer system, and wherein the second computer system is configured to calculate the second corresponding load from the one or more load factors, and wherein the first computer system and the second computer system are configured to exchange the first corresponding load and the second corresponding load to select the first virtual machine for migration.

15

9. The cluster as recited in claim 1 wherein the first virtual machine has a corresponding load that is calculated as a weighted combination of measurements of usage of two or more resources of the first computer system.

20 10. The cluster as recited in claim 9 wherein the measurements of usage include an amount of time that the first virtual machine is executing in a central processing unit of the first computer system.

25 11. The cluster as recited in claim 9 wherein the measurements of usage include an amount of input/output activity generated by the first virtual machine during execution.

12. The cluster as recited in claim 9 wherein the measurements of usage include an amount of memory occupied by the first virtual machine.

13. The cluster as recited in claim 1 wherein each of the plurality of computer systems include a schedule having a plurality of entries, each entry corresponding to a virtual machine to be executed on the respective one of the plurality of computer systems, and wherein migrating the first virtual machine comprises deleting the entry corresponding to 5 the first virtual machine in the schedule of the first computer system and inserting the entry corresponding to the first virtual machine in the schedule of the second computer system.

14. A method comprising:

10

scheduling one or more virtual machines for execution on hardware comprising a first computer system of a plurality of computer systems;

15

the first computer system detecting that the first computer system has a first load that exceeds a second load of a second computer system of the plurality of computer systems; and

20

the first computer system migrating at least a first virtual machine executing on the first computer system to a second computer system of the plurality of computer systems responsive to the detecting.

15. The method as recited in claim 14 further comprising executing the first virtual machine on the second computer system independent of the first computer system, even if the first virtual machine was initially launched on the first computer system.

25

16. The method as recited in claim 14 further comprising:

selecting the second computer system to compare loads; and

selecting the first virtual machine to migrate to the second computer system responsive to the first load exceeding the second load.

17. The method as recited in claim 16 wherein selecting the second computer system is
5 random.

18. The method as recited in claim 17 further comprising each of the plurality of computer systems periodically randomly selecting another one of the plurality of computer systems to compare loads and to potentially migrate virtual machines.

10

19. The method as recited in claim 14 wherein the first virtual machine has a corresponding load that is nearest, among the virtual machines executing on the first computer system, to 1/2 the difference between the first load and the second load.

15 20. The method as recited in claim 14 wherein the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load, the method further comprising:

20 the first computer system transmitting one or more load factors to the second computer system;

the second computer system calculating the second corresponding load from the one or more load factors; and

25

the first computer system and the second computer system exchanging the first corresponding load and the second corresponding load to select the first virtual machine for migration.

21. The method as recited in claim 14 further comprising calculating the load of the first virtual machine as a weighted combination of measurements of usage of two or more resources of the first computer system.

5 22. The method as recited in claim 14 wherein each of the plurality of computer systems include a schedule having a plurality of entries, each entry corresponding to a virtual machine to be executed on the respective one of the plurality of computer systems, and wherein migrating the first virtual machine comprises:

10 deleting the entry corresponding to the first virtual machine in the schedule of the first computer system; and

inserting the entry corresponding to the first virtual machine in the schedule of the second computer system.

15

23. A computer accessible medium encoded with a plurality of instructions that, when executed on a first computer system:

20 select a first virtual machine from one or more virtual machines to be scheduled for execution on the first computer system responsive to a first load of the first computer system exceeding a second load of a second computer system of a plurality of computer systems including the first computer system; and

25 migrate the first virtual machine to the second computer system to be executed on the second computer system.

24. The computer accessible medium as recited in claim 23 wherein the plurality of instructions, when executed, schedule the one or more virtual machines for execution on

hardware comprising the first computer system.

25. The computer accessible medium as recited in claim 22 wherein the first virtual machine executes on the second computer system independent of the first computer system during use, even if the first virtual machine was initially launched on the first computer system.
5
26. The computer accessible medium as recited in claim 23 wherein the plurality of instructions, when executed, select the second computer system to compare loads.
10
27. The computer accessible medium as recited in claim 26 wherein the second computer system is randomly selected from the plurality of computer systems.
28. The computer accessible medium as recited in claim 23 wherein the first virtual machine has a corresponding load that is approximately 1/2 the difference between the first load and the second load.
15
29. The computer accessible medium as recited in claim 28 wherein the corresponding load of the first virtual machine represents the actual load experienced in executing the first virtual machine on the first computer system.
20
30. The computer accessible medium as recited in claim 28 wherein the corresponding load of the first virtual machine represents the target load programmed for the first virtual machine in the first computer system.
25
31. The computer accessible medium as recited in claim 23 wherein the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load, and wherein the plurality of instructions,

when executed:

transmit one or more load factors to the second computer system, wherein the
second computer system is configured to calculate the second
corresponding load from the one or more load factors; and

5

exchanges the first corresponding load and the second corresponding load with the
second computer system to select the first virtual machine for migration.

10 32. The computer accessible medium as recited in claim 23 wherein the plurality of
instructions, when executed, calculate the load of the first virtual machine as a weighted
combination of measurements of usage of two or more resources of the first computer
system.

15 33. The computer accessible medium as recited in claim 23 wherein each of the plurality
of computer systems include a schedule having a plurality of entries, each entry
corresponding to a virtual machine to be executed on the respective one of the plurality of
computer systems, and wherein the plurality of instructions migrate the first virtual
machine by:

20

deleting the entry corresponding to the first virtual machine in the schedule of the
first computer system; and

25

inserting the entry corresponding to the first virtual machine in the schedule of the
second computer system.